sPHENIX activities

Chris Pinkenburg

MDC1: Simulation → Reconstruction chain

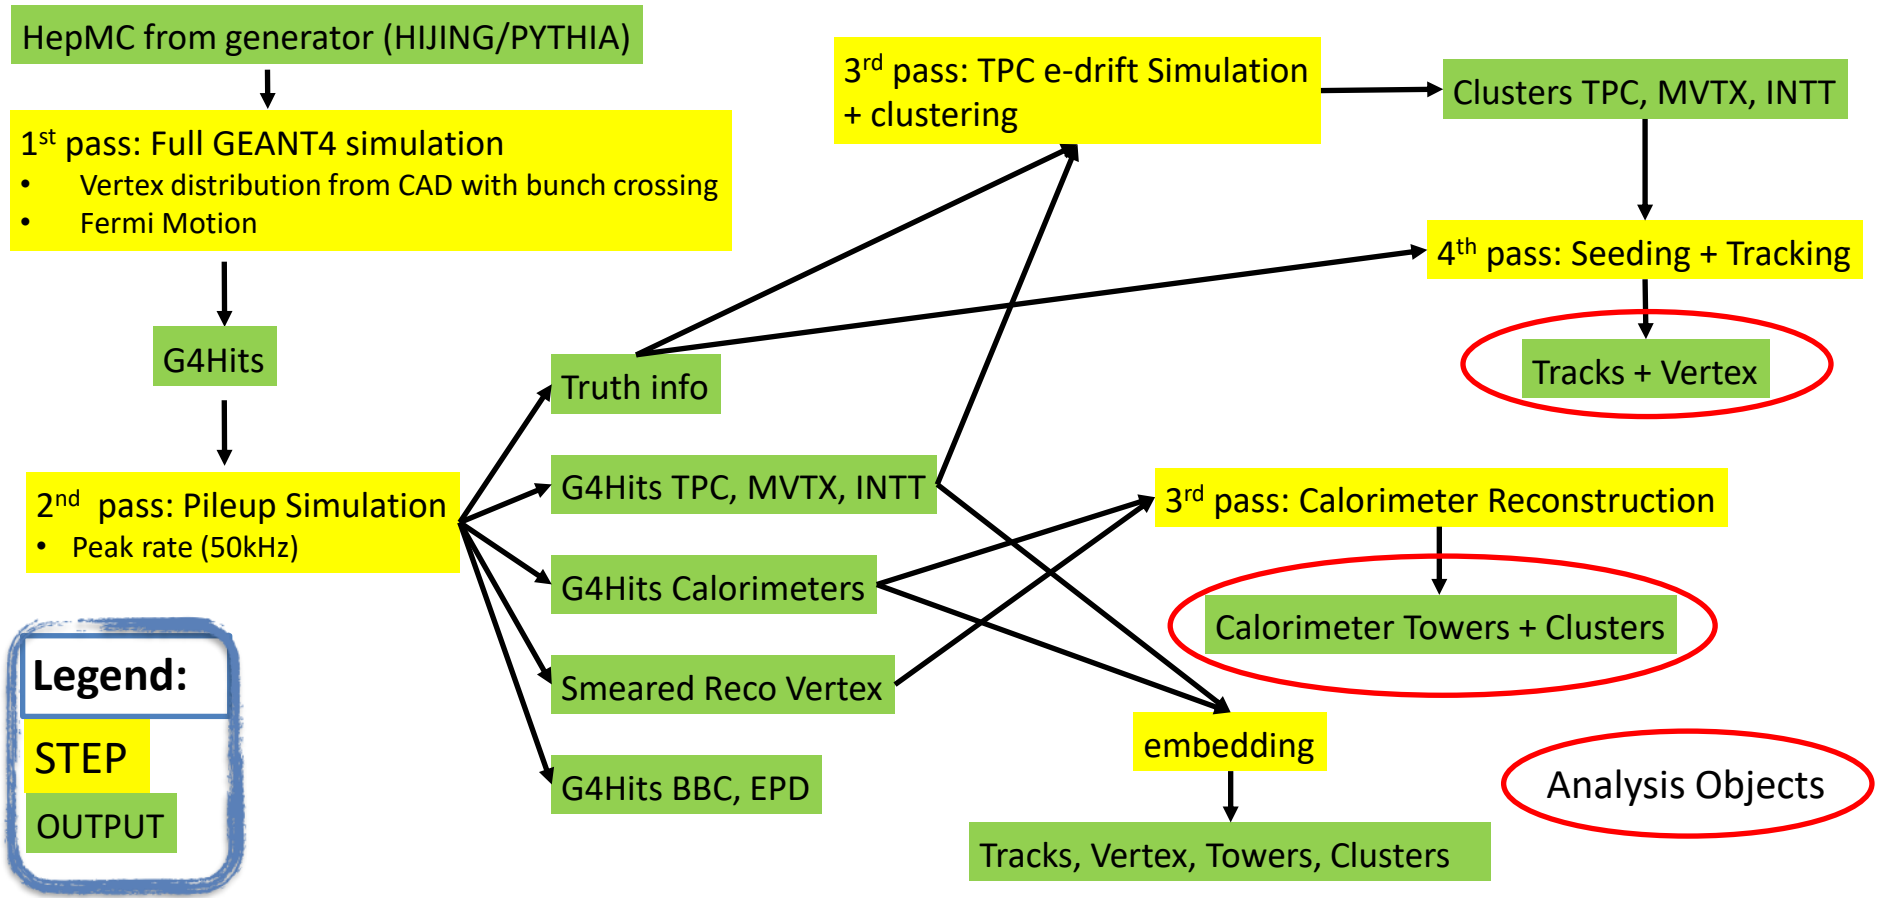

Mock Data Challenge MDC1

- Processing done in multiple steps
  - Steps typically along logical boundaries can easily be reconfigured
  - Steps can easily be rerun, chain picks up from there
- Total Running time (ballpark)
  - Heavy Ion: 48 hours/100 event MB, 50 event central
  - pp: 24 hours/1000 events
- Total data volume so far 2.3PB
- Currently running central AuAu to reach 2021 goal
  - Another pp production (different sample than existing pp) being prepared

Simulation needs

- 2021 goal
  - 500M pp
  - 6M central AuAu
- 2022 goal
  - 500M pp
  - 500M pAu
  - 10M central AuAu
  - 20M MinBias AuAu
  - 500M pp embedded

- Achieved since Nov 2020
  - 100M pp
  - 3M AuAu
  - <1M AuAu/month
  - ~100 M pp/month
  - pp impacts AuAu

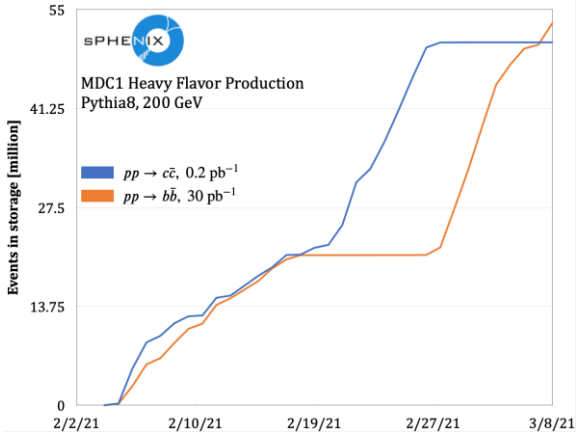

At the current pace 2021 simulations will be finished and available 2022, necessary improvements to the simulations cannot be done

Current Simulation Jobs

- Distinct classes of jobs, memory taken from condor logs
  - pp ~3GB, all steps in one
  - AuAu
    - G4: 20GB
    - TPC edrift 12GB
    - Calo clustering 2.5GB

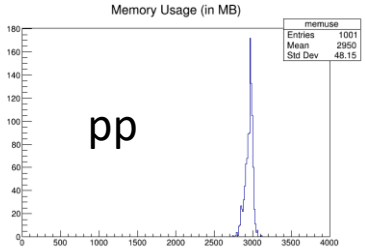

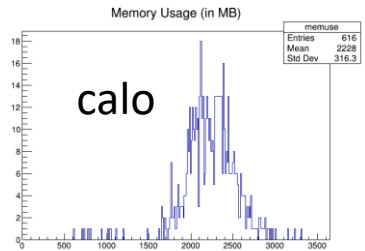

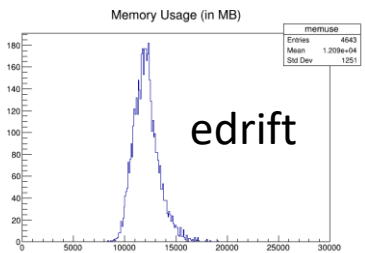

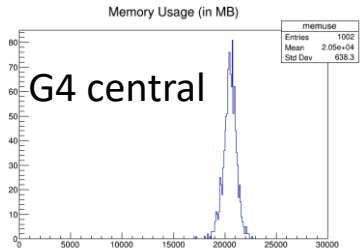

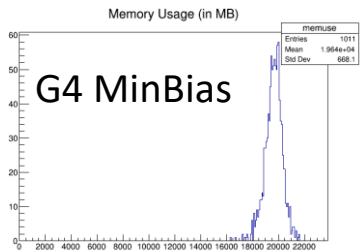

In Heavy Ions most of the memory is consumed by event data, using multi threading for processing multiple events in parallel does not help

No apparent difference between MinBias and central, condor logs contain maximum memory allocation?

The challenge

- The occasional high multiplicity event requires a large amount of memory
  - That memory drives the condor job requirement to keep eviction rate low
  - For every 2GB another core has to idle before the job gets scheduled (20GB == 10 cores)
  - Low chance of that if the farm is busy, single core jobs push our jobs out – normally we do not even run the number of jobs matching our quota
  - Adding all of PHENIX (20,000 slots in total) did not increase the number of running sPHENIX simulations, we hover at 300-700 jobs, no matter what

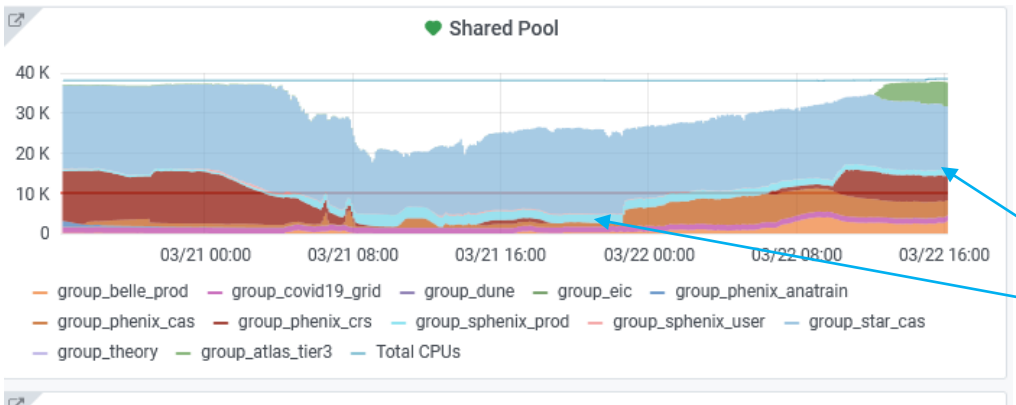

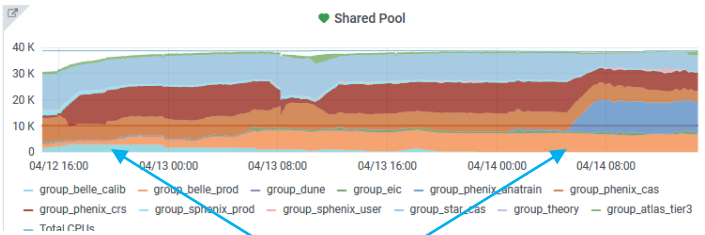

sPHENIX MDC jobs

Summary

- With the review in the rearview mirror we concentrate on continuous simulation production running
- I don’t see our code crashing – no waste of resources
- Simulations need to be refined (→ MDC2)
- Some ideas how to reduce memory but HI unrealistic to get simulation to < 4GB
- With pp@3GB, HI simulations should co-exist peacefully on 4GB nodes
- How can condor for sPHENIX be optimized in the meantime so we get at least our allocation of resources?
- dCache is working well, no user complaints – but I do see sometimes many hours of wait before a file gets copied locally